A Hybrid System for Diagnosing Multiple Disorders

Yeona Jang
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
Yeona@medg.lcs.mit.edu

Abstract

This paper investigates the advantages of introducing feedback between the processes of automated medical diagnosis and automated diagnostic-knowledge acquisition. Experimental results show that a diagnostic system with such feedback is capable of an efficiency/accuracy tradeoff when applied to the problem of diagnosing multiple disorders.

A primary feature of this work is a new mechanism, called the "diagnostic-unit" representation, for remembering results of previous diagnoses. The diagnostic-unit representation is explicitly tailored to capture the most likely relationships between disorders and clusters of findings. Unlike typical bipartite "If-Then" representations, the diagnostic-unit representation uses a general graph representation to efficiently represent complex causal relationships between disorders and clusters of findings.

In addition to the basic diagnostic-unit concept, this paper presents experience-based strategies for incrementally deriving and updating diagnostic units and the various relationships between them. Techniques for selecting diagnostic units relevant to a given problem and then combining them to generate solutions are also described.

1 Introduction

This paper addresses three interdependent problems; 1) the diagnosis problem, in particular, the problem of efficiently identifying the most likely causal events for a given body of evidence, 2) the knowledge acquisition problem, in particular, the problem of acquiring knowledge about the context sensitivity of findings that facilitates the recognition of finding patterns relevant to a problem, and 3) the problem of representing such domain-structure knowledge.

Contributions: The contribution of this work is twofold. One contribution is the development of new methods for solving problems in complex diagnostic domains. Classical techniques for diagnosis include association-based and causal-model-based reasoning. In general, an association-based reasoning system can solve problems efficiently, but is fragile in the sense that it is only good at solving familiar, prespecified problems [1, 10]. Conversely, a causal-model-based reasoning system can solve not only familiar but also unfamiliar problems from first principles, but is slow [7, 9, 12]. This work is motivated by the desire to develop a diagnostic technique whose efficiency is comparable to that of association-based reasoning but with robustness which approaches that of causal-model-based reasoning. This work shows that such a system *is* in fact possible, by combining association-based and causal-model-based reasoning.

Another contribution would be in automating the acquisition of knowledge about domain structure. It helps enhance one's understanding of the structure inherent in the diagnosis domain, by automatically discovering knowledge that identifies the structure of a problem.

Basic approach: The basic approach is the introduction of feedback between the processes of problem solving and knowledge acquisition. The introduction of such feedback results in a hybrid system that generates hypotheses to account for a given body of evidence, analyzes the results of the diagnosis, and incorporates their key features into an experiential knowledge base which, in turn, assists future diagnosis.

This research uses experience to automatically acquire domain-structure knowledge which allows decompositional abductive diagnosis to be done efficiently and effectively. For an intuitive understanding of decompositional abductive diagnosis, suppose that findings f_1, f_2, ..., f_n are given for diagnosis. One reasonable question to ask is "Can we group these findings into relatively independent subsets of findings for which the most likely explanations can be immediately inferred?" In other words, can we solve this problem by decomposition and abduction? If so, a diagnostic solution to the original problem can be generated quickly, by combining partial solutions to the subsets of findings, where each partial solution represents a disorder that explains part of the overall malfunction.

Problem decomposition allows a complex problem to be solved efficiently, by simplifying it into subproblems. Abduction also allows efficient problem solving, by avoiding reasoning step-by-step from first principles. It is the potential efficiency that motivates the use of decomposition and abduction for multidisorder diagnosis. Decompositional techniques are efficient, however, only when a problem is decomposed correctly. Unfortunately, the task of finding a correct decomposition for diagnosis is a

This research was supported in part by the National Heart, Lung, and Blood Institute under the grant number R01 HL33041, in part by the National Institutes of Health under the grant number R01 LM04493 from the National Library of Medicine, and in part by the Harvard, MIT, and Tufts/New England Medical Center Combined Research Training Program in Medical Informatics.

difficult task, for there are exponentially many ways of decomposing a given set of findings. Similarly, abduction is an effective technique for solving the diagnosis problem only when all known causal relations are most likely. In light of these observations, this research attempts to explicitly represent ***knowledge about the context sensitivity of the conclusions that can be drawn from findings.*** Such knowledge provides guides for decomposing a given set of findings and also can be formulated to capture only the most likely causal relations.

This paper describes techniques designed to address the issues of how to represent, and efficiently use, knowledge about the context sensitivity of findings, and of how to automatically acquire, from experience, such knowledge in the representation that will be described in Section 3.

Related Work: What to learn is not independent of the goals of problem solving. In the light of the interrelationship between problem solving and learning, systems such as SOAR coupled learning to problem solving. SOAR is a rule-based general-purpose problem-solving system which is integrated with explanation-based learning [8].

Another general purpose problem-solving system that incorporates several learning mechanisms is PRODIGY [11]. Much of learning in PRODIGY is directed at automatically acquiring control rules from experience to improve efficiency of a search process.

CASEY [6] is a diagnostic system that stores each solved case individually as an independent atom for use in later diagnosis. It bases diagnosis on comparison-based case-based reasoning which directly inspects old cases and transfers an entire solution of the best matching old case to a new case.

1.1 Diagnostic Problems

Input: Diagnostic problems are represented in the form of a set of findings. Findings include the history of a patient, subjective symptoms such as a patient's complaints, physical examination, and objective signs revealed either by observations or by various special laboratory tests.

Output: A *diagnosis* consists of not only the disorders primarily suspected of causing a given set of findings, but also an underlying pathophysiologic mechanism that explains how these disorders are producing all of the findings in the set. It is important to provide a underlying pathophysiologic mechanism, particularly when intermediate links within the causal chain are important determinants in the appropriate therapy for the patient. In addition, without a pathophysiologic mechanism that explains how findings are related to their primary suspected disorders, it is difficult to determine whether the results of a diagnosis "make sense" or are little more than random guesses.

This research uses a general graph notation to represent causal explanations for findings. An example of this representation is shown in Figure 1. Each black rectangular node represents an elemental disorder. Elemental disorders are either pathophysiologic states defined at a level needed for differential diagnosis or pathophysiologic states that do not require any further causes [9]. Oval nodes represent intermediate states. Intermediate states are the rest of pathophysiologic states. Rectangular nodes represent findings. Links are causal links that represent direct causal relations between the clinical entities represented by the corresponding nodes. The causal graph shown in Figure 1 represents a causal explanation with the elemental disorders producing the findings via the intermediate states identified in the graph. The same set of findings can be explained in many ways. The main driving force of this research is the question "How can we efficiently find causal graphs that represent the most likely causal explanations for a set of findings?"

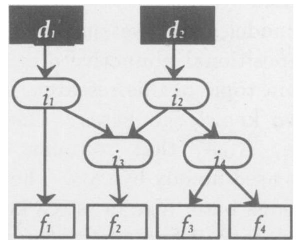

Figure 1: An example of a causal explanation

1.2 Guide to the Paper

The remainder of this paper is organized as follows. Section 3 describes methods designed to transform experiences (observations about diagnoses) to knowledge (general problem-solving rules for decompositional abductive diagnosis), while Section 4 describes techniques for using such experiential knowledge to solve diagnostic problems. The efficiency and effectiveness of these techniques are tested in the domain of heart failure diagnosis, by implementing a computer system called HYDI. An evaluation of HYDI's performance is presented in Section 5.

2 HYDI

To facilitate discussion, this section presents an overview of the system HYDI. HYDI diagnoses multiple disorders, and automatically acquires knowledge about the context sensitivity of findings from its own problem-solving experience. It performs diagnosis based on the multifault assumption without assuming that disorders are independent. HYDI also combines causal-model-based reasoning and association-based reasoning in an effective fashion.

Architecture of HYDI: The problem-solving component of HYDI consists of the causal-model-based problem solver, CMS, and the association-based problem solver, AS. The intention is to take advantage of the robustness of causal-model-based reasoning and the efficiency of association-based reasoning. Since the focus of this research is not on the specifics of causal-model-based reasoning, this research reuses an existing probabilistic causal-model-based system, HF, as CMS. HF [9] is a probabilistic causal-model-based system for heart failure diagnosis. A decompositional abductive diagnosis approach, which is the main topic of this research, is used for AS.

HYDI has two knowledge bases. One is the causal knowledge base, K_{CMS}, that contains domain causal knowledge. It is used mainly by CMS. The other is the associative knowledge base, K_{AS}, in which knowledge about the context sensitivity of findings is stored for use by AS. HYDI's problem-solving "flow" is as follows: Given a diagnostic problem, AS first tries to generate acceptable solutions to the problem. If it fails, then the more robust CMS solves the problem.

Whenever a diagnostic problem is solved, the results of the diagnosis (specifically, the most likely causal explanations for the given set of findings) are analyzed and incorporated into K_{AS} by a knowledge incorporation process, for AS's use in future diagnosis.

3 Transformation of Experience to Knowledge

This research uses experience to acquire diagnostic control knowledge for guiding decompositional abductive diagnosis. The issue of how to remember experience for use in later diagnosis must be addressed. More specifically, one needs to decide whether to treat each solved case as "atomic" or decomposable. A typical approach taken in most existing case-based systems [2, 5, 13] is to store each solved case individually as an independent atom. While easy to implement, storing cases in this way can limit the reusability of previous cases in future problem solving, by reducing the possibility of finding matches – particularly when only parts of cases match. It may also result in inefficient use of memory space, because even very similar cases are stored separately. In an attempt to deal with these difficulties, this research attempts to store solved cases in a "decomposed-and-merged" form, by analyzing results of diagnoses with respect to disorders.

3.1 Diagnostic-Unit Representation

The *diagnostic-unit representation* is a new mechanism that forms a basis for storing components of decomposed cases. In most medical domains, the significance of a finding depends on other findings that occur together with it. As a consequence, changing some findings may even require findings which remain the same to be explained differently. The diagnostic-unit representation seeks to capture this context-sensitivity of findings, by explicitly representing disorders and sets of findings that are in the most likely causal relation as units. Unlike typical bipartite "If-Then" representations, the diagnostic-unit representation uses a general graph representation to efficiently represent more complex causal relationships between disorders and sets of findings. The diagnostic-unit representation consists of two building blocks: diagnostic units and links between them.

Diagnostic units: A *diagnostic unit* is a causal graph, with a single elemental disorder root, such that previous experiences "indicate" that the causal explanation identified by the graph can be immediately inferred to be the most likely causal explanation for the findings in the graph. For example, Figure 2 shows a diagnostic unit rooted at chronic mitral regurgitation. The diagnostic unit represents that experiences so far indicate that chronic mitral regurgitation and the underlying causal mechanism identified in the graph are believed to be the most likely causal explanation for the findings in the graph.

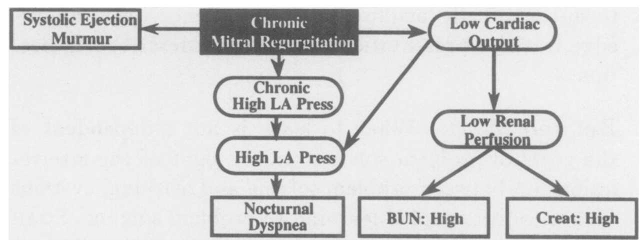

Figure 2: An example of a diagnostic unit

Diagnostic units are designed to identify diagnostic contexts in which approximate optimality is implied. A diagnostic context represents that when all of the findings identified in a diagnostic unit occur together, the disorder and underlying causal mechanism identified in the diagnostic unit can be concluded as a diagnosis for the findings. A diagnostic unit is approximately optimal if, according to past experience, the disorder and underlying mechanism identified in the diagnostic unit are believed to be the most likely diagnosis for the findings.

Diagnostic-unit sets: It is not uncommon in medical domains that different sets of findings can suggest the same elemental disorder via different underlying pathophysiologic mechanisms, thus corresponding to different diagnostic units. In light of the possibility that an elemental disorder can be associated with more than one diagnostic unit, a diagnostic-unit set for an elemental disorder is defined.

Links between diagnostic units: Two types of links are defined to represent relationships between diagnostic units. One type of link is called a *causal-relation link.* A causal-relation link represents a causal dependency between diagnostic units. The other type of link is called a *non-causal relation link,* and represents a dependency between diagnostic units which are not causally related but share common nodes.

DOC Graphs: Knowledge represented in the diagnostic-unit representation can be conceptualized as a graph where each node represents a diagnostic-unit set and each link represents a relationship between diagnostic units in different diagnostic-unit sets. Such a graph is called a *diagnostically-operative causal graph*, or DOC graph.

Detailed formal definitions of diagnostic units, of links between diagnostic units, and of DOC graphs can be found in [3].

3.2 Knowledge Incorporation Process

This subsection addresses the issue of how to acquire diagnostic units and various relationships between them. One approach would be manual compilation by interviewing expert physicians. It is, however, expensive and time-consuming. In an attempt to deal with difficulties of manual compilation, this research seeks to automatically acquire diagnostic units and relationships between them, by analyzing HYDI's own diagnostic experience.

The knowledge incorporation process assimilates new diagnostic solutions into the existing associative knowledge base, K_{AS}, of HYDI. This research assumes that K_{AS} is initially empty. In other words, no diagnostic units are known *a priori*. As diagnostic problems are solved, diagnostic units and relationships between them are discovered and updated in K_{AS} for use in later diagnosis. This is done by the knowledge incorporation process that consists of the "DOC transformation process" followed by the "joining-up process."

3.2.1 DOC Transformation Process

The DOC transformation process constructs DOC graphs for new diagnostic solutions. Given a new solution, it first compiles diagnostic units from a new diagnostic solution by collecting, for each elemental disorder in the diagnostic solution, all nodes and causal links in the diagnostic solution that are reachable from the elemental disorder. For example, suppose that a diagnostic problem is solved, and the causal graph in Figure 1 is generated as the mostly likely causal explanation for the problem. Then, the two diagnostic units rooted at d_1 and d_2, respectively, can be compiled from this diagnostic solution. While there is no causal path between d_1 and d_2 in the diagnostic solution, the two diagnostic units share a common node i_3, allowing a non-causal relation link between them to be established. Figure 3 shows the DOC graph of the diagnostic solution in Figure 1.

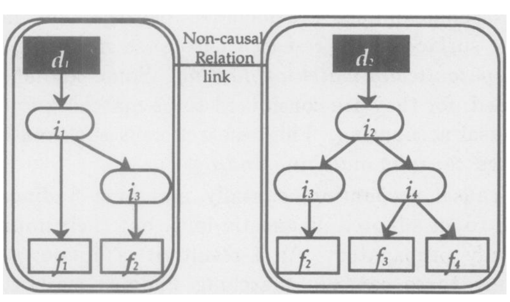

Figure 3: A DOC graph for the solution in Figure 1

3.2.2 Joining-Up Process

The joining-up process incorporates newly acquired DOC graphs into K_{AS}. A fundamental issue that must be addressed by the joining-up process is whether a new diagnostic unit should be used to update an existing diagnostic unit in K_{AS} or be considered to be a new element of the corresponding diagnostic-unit set. This merging issue is addressed with a *merging threshold*. If two diagnostic units are a "partial recovery" of the same diagnostic unit, then they are likely to have a similar underlying structure. Two diagnostic units are merged to produce a combined whole if they have the roots representing the same elemental disorder and the underlying structure of the two diagnostic units match by more than certain percentage, *i.e.*, the merging threshold.

The joining-up process also performs bookkeeping for use in later diagnosis, updating statistics such as how frequently a diagnostic unit has occurred so far. In addition, dependencies between diagnostic units are also updated in K_{AS}.

4 Decompositional Abductive Diagnosis

The previous section addressed the problem of how to acquire and represent knowledge about the context sensitivity of findings associated with particular disorders. This section addresses the issue of how to use such experiential knowledge for efficient and effective decompositional abductive diagnosis.

Decompositional abductive diagnosis can be viewed as a two-stage process. The first stage is the grouping of a given body of evidence into subproblems. The basic approach is to find relevant diagnostic units based on an approximate technique called "deep matching adaptation." The second stage is to construct a diagnostic solution to the problem, from the relevant diagnostic units selected by the preceding evidence-grouping process.

4.1 Evidence-Grouping Process

The macro-finding captured in each diagnostic unit, as a whole, represents a clinical indicator that the specified set of findings strongly supports the existence of the disorder and underlying pathophysiologic mechanism identified in the diagnostic unit. In other words, diagnostic units provide guides for decomposing a set of findings into smaller, immediately solvable subsets. In the diagnostic-unit representation paradigm, therefore, the grouping of given evidence becomes a search for relevant diagnostic units. This research addresses the issue of how to determine relevant diagnostic units, by matching diagnostic units against the given evidence.

A common matching method is to simply count the number of findings that match on the surface. While easy to implement, this matching method, called "simple matching" in this research, is only effective when cases that are similar on the surface occur frequently. Unfortunately, in such medical domains as heart failure, patients with the exact same findings rarely occur. Since diagnostic units used in this paper are acquired from experiences, knowledge captured in diagnostic unit is generally incomplete. Simple matching thus appears to be less suitable for use in medical diagnosis. The issue is how to translate or adapt existing diagnostic units so that they match a new problem. To address this issue, this research investigates deep matching adaptation. A key feature of deep matching adaptation is that it considers not only similarity in appearance but also similarity in underlying causality.

The issue that arises in matching a diagnostic unit against a diagnostic problem is what to do with findings that are unmatched on the surface. There are two kinds of unmatched findings. One is unmatched findings in a diagnostic unit, and the other is unmatched findings in a diagnostic problem. Since unmatched unit findings are not known to a patient, they have no effect on the patient's state. In light of this observation, this thesis handles unmatched unit findings by removing them from the diagnostic unit. Such removal can invalidate the diagnostic unit, however. The issue regarding the validity of a diagnostic unit is addressed in the succeeding hypothesis-construction process. For unmatched problem findings, this thesis checks to see if they can be explained by the diagnostic unit. To this end, a technique called "causal accounting" is investigated.

Causal accounting: Causal accounting is a simple method for tailoring diagnostic units based on underlying causality of findings. It allows an unmatched finding in a diagnostic problem to be added to a diagnostic unit (and treated as a matching finding), if there exists in the diagnostic unit a pathophysiologic state which can directly cause the unmatched problem finding: Such a pathophysiologic state in the diagnostic unit is referred to as an "accounting state" of the unmatched problem finding. For example, suppose that the diagnostic unit in Figure 2 is being matched against a diagnostic problem which consists of systolic ejection murmur, high BUN, high creatinine, and orthopnea. According to HF's causal model, high LA (Left Atrial) pressure can directly cause orthopnea. Causal accounting thus allows orthopnea, which is a unmatched problem finding, to be added to the diagnostic unit as an effect of high LA pressure.

Causal accounting can efficiently increase the usability of diagnostic units in problem solving. The depth of causality examined by causal accounting is limited to direct causal dependencies, however, in order to avoid a costly complete propagation of new evidence impact throughout the entire network. In this regard, causal accounting can be considered as a one-step lookahead version of a more general accounting principle. By doing so, causal accounting trades accuracy for computational efficiency.

Deep matching adaptation: The evidence-grouping process uses deep matching adaptation to find relevant diagnostic units. Deep matching adaptation tailors a diagnostic unit to a particular case in two ways – by applying causal accounting to add unmatched problem findings to the diagnostic unit and by removing unmatched unit findings from the diagnostic unit. The algorithm for deep matching adaptation can be summarized as follows: For a diagnostic unit, 1) make a copy of the diagnostic unit, 2) for each unmatched finding f in the diagnostic problem, if f can be explained by causal accounting, then add to the copy a node n that represents f and a direct causal link to n from the corresponding accounting state, 3) remove all unmatched findings in the diagnostic unit from the copy, and then 4) return the modified copy.

The evidence-grouping process applies deep matching adaptation to each diagnostic unit that exists in K_{AS}.

4.2 Hypothesis-Construction Process

The hypothesis-construction process generates a solution to the original problem by combining the adapted diagnostic units chosen by the evidence-grouping process. Note that each diagnostic unit returned by the evidence-grouping process is adapted to explain some subset of the given finding set. As a consequence, any combination of adapted diagnostic units can be an explanation for all of the given findings if the union of adapted diagnostic units in the combination is equal to the given finding set. The issue is how to find a combination which results in the most likely causal explanation. The problem is that not all adapted diagnostic unit selected by the evidence-grouping process are ones that are part of a correct diagnosis. Some of the adapted diagnostic units are, in fact, falsely chosen as relevant ones. This research calls a diagnostic unit which is part of a correct diagnosis a *true positive unit*, and a diagnostic unit which is not a part of correct diagnosis but falsely chosen a *false positive unit*. Unfortunately, during testing, the evidence-grouping process returned many diagnostic units, and most of them were false positive: More specifically, the evidence-grouping process returned an average of 66 adapted diagnostic units, and over 90 how to pick "correctly" true positive units from the output of the evidence-grouping process.

Specificity-Reflected Similarity Metric: A common picking method is to count the number of matching findings, and then pick the units with the largest number. This method, called a "simple similarity metric" in this research, is easy to implement. Not all matching findings in an adapted diagnostic unit are, however, of the same kind. Some findings are included because they match on the surface. This research calls such matching findings *syntactically matching finding*. Some findings are included, for they are considered to be matching findings by causal accounting. This research calls such matching findings *causally matching findings*.

In causal accounting, causally matching findings are added to an adapted diagnostic unit, but their impact is not fully propagated. As a result, it is unknown how strongly these causally matching findings support the existence of the diagnostic unit, while for syntactically matching findings, it is indicated by previous experience that they strongly do. This research investigates a similarity metric, called *specificity-reflected similarity metric*, to take this difference into account.

The specificity-reflected similarity metric is motivated by the observation that some findings do better than others in identifying the existence of a disorder. Findings in a diagnostic unit are divided into two groups: specific findings and non-specific findings. Specific findings of a diagnostic unit are findings that play a significant role in identifying the existence of the diagnostic unit, and non-specific findings are the remaining findings. The "playing a significant role" clause is implemented via a comparison of the specificity attached to each finding in a diagnostic unit with a threshold. The specificity of a finding represents a level of the significance of a finding in identifying a particular disorder and the underlying pathophysiologic mechanism. HF is capable of providing specificities that range between 0 and 1, which the knowledge incorporation process of HYDI remembers for each finding in a diagnostic unit. The remembered specificity of a finding in a diagnostic unit is compared with a threshold, to determine if the finding can be considered to be specific. During testing, a finding in a diagnostic unit with the specificity higher than 0.8 was considered as a specific finding of the diagnostic unit.[1]

Unlike the simple similarity metric, the specificity-reflected similarity metric uses three measures: the number of matching specific findings, the total number of matching findings, and frequency with which a diagnostic unit has occurred so far. The specificity-reflected similarity metric determines ranks of adapted diagnostic units as follows: For any two adapted diagnostic units, the one with the larger number of matching specific findings gets higher rank. In case of a tie, the one with the larger total number of matching findings gets higher rank. In case of a tie, the one that has occurred more commonly gets higher rank.

Dependency-Guided Picking Method: If the adapted diagnostic unit picked by the specificity-reflected similarity metric does not explain all of the given findings, other adapted diagnostic units need to be chosen to account for the unexplained findings. A common approach is to choose these units based on the disorder independence assumption. In most medical domains, however, disorders are not always independent of each other. In light of such dependency between disorders, this research uses dependencies between diagnostic units to guide the picking process. Basically, additional units are chosen by applying the specificity-reflected similarity metric to adapted diagnostic units which have causal or non-causal relation links to the adapted diagnostic units which have already been picked. If there are no such adapted diagnostic units, all remaining adapted diagnostic units are considered. This picking process is repeated until all of the given findings are explained.

5 Empirical Analysis

This section presents an experiment conducted to test the efficiency and effectiveness of the techniques developed in this paper, by implementing HYDI.

The data set used for the empirical analysis consisted of 300 cardiac patients, from The New England Medical Center Hospital. In order to reduce bias due to case ordering, 50 independent trials were conducted on 50 different random case orderings, and their results were averaged together. In each trial, no diagnostic units were known in advance, and all 300 patients were run. Each time a diagnostic problem was solved, a diagnostic solution was incorporated into K_{AS} which was, in turn, used in the next diagnosis.

Diagnostic performance was measured along the dimensions of accuracy and running time. In particular, accuracy was measured in terms of true-positive accountability and false-positive rate. True-positive accountability is the percentage of given findings explained by true positives in a diagnostic solution. False-positive rate is the percentage of diagnostic units in a diagnostic solution that are false positives. Running time was measured on a SUN SPARC Station 2.

5.1 Diagnostic Performance of HYDI

HYDI used HF as its causal-model-based solver. For HYDI's association-based solver, decompositional abductive diagnosis (for expository convenience, called DAD) was implemented that selects relevant diagnostic units based on deep matching adaptation, and constructs a diagnostic solution by applying the specificity-reflected similarity metric and the dependency-guided picking method. Detailed algorithms for DAD can be found in [4].

HYDI performs hybrid reasoning to solve diagnostic problems. The central idea of hybrid reasoning is that for a diagnostic problem, DAD is first performed to solve the problem; if a diagnostic solution generated by DAD is not acceptable, then HF is called to solve the problem. Two strategies for determining whether or not a diagnostic solution generated by DAD is acceptable were considered. One strategy is to accept a causal explanation generated by DAD if the explanation can account for all the findings in the problem. This strategy was implemented in HYDI_1. The other strategy is to accept a causal explanation if the explanation not only can explain all the findings in the problem, but also is "close to a correct diagnosis." The "close to a correct diagnosis" clause was implemented as a comparison of true-positive accountability with a threshold. During testing, 88%, which is the average true-positive accountability achieved by HYDI_1, was used as the threshold value. This strategy was implemented in HYDI.

Table 1 summarizes average percentage of findings in a diagnostic problem that are explained by a diagnostic solution. Unlike the three other systems, DAD was able to explain only an average of 89% of findings in a diagnostic problem. This empirically verifies that the association-based DAD is less robust than the other systems.

Table 2 summarizes the results of the experiment, in terms of average accuracy and running time. During evaluation, the most likely causal explanations generated by HF were assumed to be "correct" diagnoses. In other words, HF has 0 false-positive rate and 100% true-

[1] In order to determine a threshold value, a test was conducted to learn the threshold which gives the best outcome in identifying diagnostic units. The test showed that the threshold of 0.8 produced the best outcome.

A diagnostic solution S	Average % of findings explained by S
generated by DAD	88.9%
generated by HF	100%
generated by HYDI$_1$	100%
generated by HYDI	100%

Table 1: Average percentage of findings explained

positive accountability. The experiment shows that a diagnostic solution generated by HF consisted of an average of 8 diagnostic units, but with only around 40% of them being ones that had been seen previously.

System	Accuracy		Avg. Running Time (σ)
	Avg. TPA	Avg. FPR	
HF	100%	0	52.3 sec (138.2 sec)
DAD	81.0%	13.1%	3.7 sec (12.2 sec)
HYDI$_1$	88.1%	10.9%	4.0 sec (13.6 sec)
HYDI	96.8%	4.3%	17.7 sec (43.1 sec)

Table 2: Average accuracy and running time

As summarized in Table 2, the experiment indicates that HYDI achieved an average of 96.8% true-positive accountability and 4.3% false-positive rate. This result may appear disappointing, but in fact is remarkably good, given the fact that most of the problem solving in HYDI was done by AS (which based its problem solving on previous experience), and the fact that similar cases occurred infrequently (only an average of 40% of the diagnostic units in a solution were seen previously). The results shown in the table also empirically suggest that attractive tradeoffs between accuracy and efficiency can be achieved. In particular, compared with HF, HYDI was able to achieve a 300% increase in speed, at only a 3% decrease in accuracy. HYDI also shows a smaller standard deviation in running time than HF.

Finally, the experiment shows that as experience was accumulated, CMS was called less frequently for diagnosis. In HYDI$_1$, after an average of 27 cases, DAD was generally able to find some diagnostic solutions that explain all the findings. In HYDI, even though CMS was called more frequently than in HYDI$_1$, it was still called less as experience was gathered.

6 Concluding Remarks

This paper described a framework, called the diagnostic-unit representation, for structuring diagnostically critical knowledge as graphs. It also described possible strategies for incrementally acquiring such graphs by analyzing the results of diagnoses with respect to disorders. In addition, this paper explored the following new techniques for decompositional abductive diagnosis; deep matching adaptation to find relevant diagnostic units, the specificity-reflected similarity metric to determine a level of relevance of a diagnostic unit, and a picking method which uses dependencies between disorders.

As empirically demonstrated, the diagnostic-unit representation appears to be an effective way of capturing domain decomposability, thereby facilitating one's understanding of the structure of a problem, and allowing decompositional abductive diagnosis to be done efficiently and effectively. Inappropriate use of knowledge can adversely affect overall diagnostic performance. Empirical results demonstrates that the techniques developed for using diagnostic units are effective, especially when similar cases rarely occur.

Currently, the knowledge incorporation component of HYDI does not use failures to learn. HYDI could be extended to address issues such as how failures can be used to change the structure of existing knowledge.

Acknowledgments

The author would like to thank Peter Szolovits for his support and numerous stimulating comments, William Long for his help in understanding the heart failure domain and the HF program, and Alexander Ishii for his draft-reading efforts and insightful suggestions.

References

[1] B.G. Buchanan and E.H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, MA, 1984.

[2] K.J. Hammond. "CHEF: A model of case-based planning." In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 267-271, 1986.

[3] Y. Jang. "Knowledge Representation and Incorporation in a Hybrid Reasoning System with Feedback," in B. Nevel, C. Rich, and W. Swartout, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR92)*. Morgan Kaufmann, San Mateo, CA, 1992.

[4] Y. Jang. "*HYDI:* A Hybrid System with Feedback for Diagnosing Multiple Disorders." Ph.D dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1993.

[5] J.L. Kolodner. "Retrieving the Best Case for a Case-Based Reasoner." In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,* pages 155-162, Lawrence Erlbaum, 1989.

[6] P.A. Koton. "Using experience in learning and problem solving." TR 441, Massachusetts Institute of Technology, Laboratory for Computer Science, Cambridge, MA, 1989.

[7] C.A. Kulikowski and S.M. Weiss. "Representation of expert knowledge for consultation: The CASNET and EXPERT projects." In P. Szolovits, editor, *Artificial Intelligence in Medicine*, Westview Press, 1982.

[8] J.E. Laird, P.S. Rosenbloom, and A. Newell. "Chunking in SOAR: The anatomy of a general learning mechanism." In *Machine Learning*, 1, pages 11-46, 1986.

[9] W.J. Long. "Medical diagnosis using a probabilistic causal network." *Applied Artificial Intelligence*, 3:367-383, 1989.

[10] R.A. Miller, H.E. Pople, Jr., and J.D. Myers. "INTERNIST-1, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine." In W.J. Clancey and E.H. Shortliffe, editors, *Readings in Medical Artificial Intelligence*, pages 190–209, Addison-Wesley, Reading, MA, 1984.

[11] S. Minton, J.G. Carbonell, C.A. Knoblock, D.R. Kuokka, O. Etzioni, and Y. Gil. "Explanation-based Learning: Optimizing Problem Solving Performance through Experience." *Paradigms for Maching Learning*, 1989.

[12] R.S. Patil. "Causal representation of patient illness for electrolyte and acid-base diagnosis." MIT TR-267, MIT Lab. for Computer Science, 1981.

[13] E. Rissland and K. Ashley. "Credit Assignment and he Problem of Competing Factors in Case-Base Reasoning." In *Proceedings of DARPA Workshop on Case-Based Reasoning,* pages 327-344, 1988.